

UNITED STATES PATENT APPLICATION

FOR

**ADDRESS CONVERSION TECHNIQUE IN A CONTEXT SWITCHING
ENVIRONMENT**

INVENTOR:

Dennis M. O'Connor

INTEL CORPORATION

Prepared by:
Rita M. Wisor
Reg. No. 41,382

BACKGROUND

Description of the Related Art

[1001] A processor in a computing system uses memory mapping to map virtual addresses to physical addresses. The address space generated by, for example, the execution unit of the processor, is referred to as the virtual address space. The addresses that correspond to hardware memory locations available on the system are referred to as physical address.

[1002] A processor may implement an additional level of address remapping for process or context switching. A process identifier associated with a context may be used to remap addresses. The use of process identifiers and remapping alleviates the need to flush or invalidate a cache on process switches.

BRIEF DESCRIPTION OF THE DRAWINGS

[1003] The present invention may be better understood, and its numerous features and advantages made apparent to those skilled in the art by referencing the accompanying drawings.

[1004] FIG. 1 illustrates various logic blocks of a computing system according to an embodiment of the present invention.

[1005] FIG. 2 illustrates a block diagram of a memory management unit according to an embodiment of the present invention.

[1006] FIG. 3 illustrates a flow diagram of virtual address to physical address conversion according to an embodiment of the present invention.

[1007] FIG. 4 illustrates a process switching flow according to an embodiment of the present invention.

[1008] The use of the same reference symbols in different drawings indicates similar or identical items.

DESCRIPTION OF THE EMBODIMENT(S)

[1009] According to some embodiments, a memory management unit receives a virtual address and provides a corresponding physical address. The memory management unit stores generated virtual address-to-physical address translations. If a virtual address-to-physical address translation is available for a particular virtual address, the memory management unit retrieves the corresponding physical address. If a translation is not available, the memory management unit generates the corresponding physical address from the virtual address. The memory management unit converts the virtual address to a modified virtual address using a process identifier and then performs a page table walk using the modified virtual address, generating the physical address.

[1010] In the following description, numerous specific details are set forth. However, it is understood that embodiments of the invention may be practiced without these

specific details. In other instances, well-known methods, structures and techniques have not been shown in detail in order not to obscure an understanding of this description.

[1011] References to “one embodiment,” “an embodiment,” “example embodiment,” “various embodiments,” etc., indicate that the embodiment(s) of the invention so described may include a particular feature, structure, or characteristic, but not every embodiment necessarily includes the particular feature, structure, or characteristic. Further, repeated use of the phrase “in one embodiment” does not necessarily refer to the same embodiment, although it may.

[1012] As used herein, unless otherwise specified the use of the ordinal adjectives “first,” “second,” “third,” etc., to describe a common object, merely indicate that different instances of like objects are being referred to, and are not intended to imply that the objects so described must be in a given sequence, either temporally, spatially, in ranking, or in any other manner.

[1013] Unless specifically stated otherwise, as apparent from the following discussions, it is appreciated that throughout the specification discussions utilizing terms such as “processing,” “computing,” “calculating,” or the like, refer to the action and/or processes of a computer or computing system, or similar electronic computing device, that manipulate and/or transform data represented as physical, such as electronic, quantities into other data similarly represented as physical quantities.

[1014] In a similar manner, the term “processor” may refer to any device or portion of a device that processes electronic data from registers and/or memory to transform that

electronic data into other electronic data that may be stored in registers and/or memory.

A “computing platform” may comprise one or more processors.

[1015] FIG. 1 illustrates various logic blocks of a computing system according to an embodiment of the present invention. A processor 100 includes an address generation unit (AGU) 102 and a memory management unit (MMU) 104. AGU 102 may be, for example, an arithmetic logic unit (ALU) that generates data addresses or, for example, an incrementor that generates instruction addresses. The addresses generated by AGU 102 are virtual addresses. MMU 104 converts the virtual addresses into physical addresses. Those physical addresses are used to access, for example, to read or write a storage 106. Storage 106 may be internal or external to processor 100.

[1016] Processor 100 represents a central processing unit of any type of architecture, including an ARM, a CISC or a RISC type architecture. Storage 106 represents one or more mechanisms for storing data. For example, storage 106 may include read only memory (ROM), random access memory (RAM), magnetic disk storage mediums, optical storage mediums, and/or flash memory devices. While one embodiment will be described in which the invention is implemented in a single processor computing system, embodiments of the invention could be implemented in a multi-processor computing system.

[1017] Processor 100 can be in any of a number of computing and communication systems including, but not limited to, mainframes, minicomputers, servers, workstations, personal computers, notepads, personal digital assistants, various wireless

communication devices that may include one or more antenna(e) 108 and transceiver 110 and embedded systems, just to name a few.

[1018] FIG. 2 illustrates a block diagram of memory management unit 104 according to an embodiment of the present invention. MMU 104 receives a virtual address and provides a corresponding physical address. First, a determination is made whether a translation for the virtual address is contained within a store, for example, a translation look-aside buffer (TLB) 202. TLB 202 contains one or more virtual address-to-physical address translations. If a translation for the virtual address is not within TLB 202, the translation is generated by TLB miss handling unit 204. TLB miss handling unit 204 includes conversion logic 206 and a page table walk unit 208. Conversion logic 206 converts a virtual address (VA) into a modified virtual address (MVA). Conversion logic 206 may include, for example, determination logic 210 to determine if one or more upper bits of the virtual address, for example, the seven upper bits VA[31:25], are equal to zero, and a multiplexer 212 to select the upper bits of a modified virtual address (MVA). For example, multiplexer 212 selects a process identifier, PID[6:0] to replace the upper seven bits of the virtual address if those bits of the virtual address are equal to zero. Otherwise, the upper seven bits of the virtual address are selected.

[1019] The process identifier (PID) may be stored in a register on processor 100 and may be associated with a particular process. For a seven-bit PID, the virtual address may be remapped to one of 128 “slots” in a 4 Gbyte address space. This process identifier remapping may be useful for operating system management of processes that map to the same virtual address space. In those cases, virtually mapped caches would not require invalidating on a process switch until, for example, such time as the PID value rolls over

or a PID value is reused. An address that has yet to be modified by the PID (“PIDified”) is referred to as a virtual address (VA). An address that has been through conversion logic 206, but not translated into a physical address, is referred to as a modified virtual address (MVA).

[1020] Page table walk unit 208 generates a physical address using the modified virtual address. Memory addressing schemes often use paging to implement virtual memory. When using paging, the virtual address space may be divided into fix-sized blocks called pages, each of which may be mapped onto any of the physical addresses. For example, page table walk unit 208 determines and maintains, according to a paging algorithm, the current mappings for the virtual to physical addresses using page tables. The page tables are often in main memory and accessing them may be time consuming. To speed up the paging translations, some of the generated virtual address-to-physical address translations are stored in TLB 202.

[1021] While these translated memory addresses are described herein as physical memory addresses, in alternative embodiments these translated memory addresses could be used for any number of purposes. For example, further translations could be performed on these translated memory addresses before physical addresses are achieved. In addition, while one embodiment is described in relation to translating a virtual memory address space, alternative embodiments could use the invention to assist in the translation of any type of addresses.

[1022] Some or all of the virtual address-to-physical address translations stored in TLB 206 may be invalidated when the process identifier is updated. The process

identifier is updated, for example, when a process or context is switched. Because TLB 202 may contain translations for virtual addresses that have null upper bits and therefore are specific to a particular process, those translations must be invalidated when the process switches. Invalidating the translations in TLB 202 may be performed by software. Alternatively, hardware may detect the change and perform the invalidation automatically.

[1023] FIG. 3 illustrates a flow diagram of processes performed by processor 100 upon a request to MMU 104 to translate a virtual address into its corresponding physical address according to one embodiment of the present invention.

[1024] MMU 104 receives a virtual address, process block 302. A determination is made whether a translation for the virtual address is stored, for example, in TLB 202, process block 304. If so, the physical address is retrieved, process block 306. If not, a determination is made whether the upper bits, for example, the upper seven bits, of the virtual address are null, that is, equal to zero, process block 308. If so, the upper bits are replaced with the process identifier (PID), process block 310. Next, using the PIDified virtual address, that is, the modified virtual address (MVA), a page table walk is performed, generating the physical address, process block 312. TLB 202 may be updated with the new virtual address-to-physical address translation, process block 314.

[1025] FIG. 4 illustrates a process switching flow according to an embodiment of the present invention. A determination is made whether the process identifier (PID) has been updated, process block 402. If the PID has been updated, all or part of the TLB is

invalidated, process block 404. The entire TLB may be invalidated. Alternatively, only the translations where the upper bits of the virtual address are null may be invalidated.

[1026] The techniques described above may be embodied in a computer-readable medium for configuring a computing system to execute the method. The computer readable media may be permanently, removably or remotely coupled to system 100 or another system. The computer readable media may include, for example and without limitation, any number of the following: magnetic storage media including disk and tape storage media; optical storage media such as compact disk media (e.g., CD-ROM, CD-R, etc.) and digital video disk storage media; holographic memory; nonvolatile memory storage media including semiconductor-based memory units such as FLASH memory, EEPROM, EPROM, ROM; ferromagnetic digital memories; volatile storage media including registers, buffers or caches, main memory, RAM, etc.; and data transmission media including permanent and intermittent computer networks, point-to-point telecommunication equipment, carrier wave transmission media, the Internet, just to name a few. Other new and various types of computer-readable media may be used to store and/or transmit the software modules discussed herein. Computing systems may be found in many forms including but not limited to mainframes, minicomputers, servers, workstations, personal computers, notepads, personal digital assistants, various wireless devices and embedded systems, just to name a few. A typical computing system includes at least one processing unit, associated memory and a number of input/output (I/O) devices. A computing system processes information according to a program and produces resultant output information via I/O devices.

[1027] Realizations in accordance with the present invention have been described in the context of particular embodiments. These embodiments are meant to be illustrative and not limiting. Many variations, modifications, additions, and improvements are possible. Accordingly, plural instances may be provided for components described herein as a single instance. Boundaries between various components, operations and data stores are somewhat arbitrary, and particular operations are illustrated in the context of specific illustrative configurations. Other allocations of functionality are envisioned and may fall within the scope of claims that follow. Finally, structures and functionality presented as discrete components in the various configurations may be implemented as a combined structure or component. These and other variations, modifications, additions, and improvements may fall within the scope of the invention as defined in the claims that follow.